

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
1 July 2004 (01.07.2004)

PCT

(10) International Publication Number
WO 2004/055692 A1

(51) International Patent Classification⁷: **G06F 17/30**

(72) Inventor; and

(21) International Application Number:
PCT/AU2003/001668

(75) Inventor/Applicant (for US only): **KOSMYNIN, Arkadi**
[AU/AU]; 2/28 Centennial Ave, Brunswick West, Victoria
3055 (AU).

(22) International Filing Date:
16 December 2003 (16.12.2003)

(81) Designated States (national): AU, US.

(25) Filing Language: English

Declarations under Rule 4.17:

(26) Publication Language: English

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii)) for all designations
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii)) for all designations
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii)) for all designations
- of inventorship (Rule 4.17(iv)) for US only

(30) Priority Data:
2002953514 16 December 2002 (16.12.2002) AU
2003902054 30 April 2003 (30.04.2003) AU

Published:

- with international search report

(71) Applicant (for all designated States except US): **OZ INSIGHT PTY LTD** [AU/AU]; 2/28 Centennial Avenue, Brunswick West, Victoria 3055 (AU).

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: A METHOD AND SYSTEM FOR INTERACTIVE WORK WITH MULTIMEDIA OBJECTS POSTED ON THE USENET

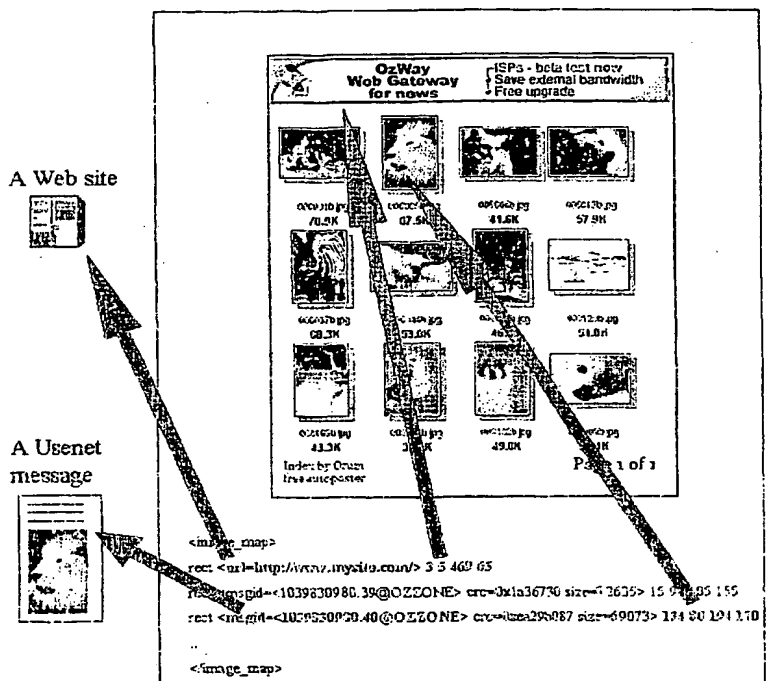


Image map records establish associations between image regions and Usenet messages or Web sites

(57) Abstract: The invented system provides for construction, requesting, obtaining and using previews of Usenet multimedia objects in order to make informed downloading decisions. Selection and downloading of Usenet images is made fast and easy when image indexes in a format disclosed in this invention are available. Each such index contains an image and an image map that is a set of records associating image regions with particular Usenet articles. This map enables unique identification of the article that user wants when they click on a particular area of the index image. The user client (newsreader) can then automatically find and start downloading the desired article. If an advertising banner is included in the image, the system can direct user to the advertised web site and allows counting visits generated by such banners. If a multimedia objects preview is not found in the newsgroup, users can send an index request to an index server. The server will download relevant articles, construct previews and either post them to the Usenet or return to the requesting user, or both.

BEST AVAILABLE COPY

A Method and System for Interactive Work with Multimedia Objects Posted on the Usenet

1 Field of Invention

The present invention relates to Internet information services. In particular, the present invention relates to improvements in use of the Usenet.

In the first aspect, the present invention relates to helping Usenet users make informed decisions on whether they want to download a particular Usenet message.

In the second aspect, the present invention relates to helping Usenet users to obtain information that is necessary to make informed decisions on whether they want to download a particular Usenet article. It describes a method and system allowing users to request small descriptions of large multimedia objects available for downloading. After previewing the description, users can make a decision on whether they want to download the objects it describes.

In the third and fourth aspects, this invention relates to advertising on the Usenet. It describes a method and system allowing inserting and using advertisement banners in Usenet messages and counting number of referred visits produced by a particular advertisement object.

2 Background

The Usenet is a worldwide bulletin board system that can be accessed through the Internet or through many online services. The Usenet contains tens of thousands of forums, called newsgroups, that cover many and varied interest groups. The Usenet is used daily by millions of people around the world.

Every Usenet message belongs to a newsgroup. Messages are made available to users worldwide by means of the UUCP and NNTP protocols (Unix to Unix Copy Program, and Network News Transport Protocol, respectively). Individual computing sites oversee the huge quantity of incoming messages, and decide how long messages can be kept before they must be removed to make room for new ones. Typically, messages are stored for less than a week. They are made available to users via a news server.

Users access local newsgroups with a newsreader program. Modern WWW browsers come with a built-in newsreader. A dedicated newsreader program can also be used. The newsreader accesses the local (or remote) News host using the NNTP, enabling a user to download as many newsgroups and their contents as they desire. If there is no local access to a news server, there are publicly accessible commercial and free Usenet hosts that can be accessed.

Users sending Usenet messages must address each message to a particular newsgroup. There are newsgroups on subjects ranging from education for the disabled to Star Trek and from environment science to politics in the former Soviet Union. The quality of the discussion in newsgroups may be excellent, but this is not guaranteed. Some newsgroups have a moderator who scans the messages for the group and decides which ones are appropriate for distribution.

Some of the newsgroups provide a useful source of information and help on technical topics. Users needing to find out about a subject often send questions to the appropriate newsgroup, and an expert somewhere in the world can often provide an answer. Lists of Frequently Asked Questions (FAQs) are compiled and made available periodically in some newsgroups.

The Usenet was originally designed for exchange of textual information, but presently the major part of bandwidth and storage resources is consumed by so called "binary" newsgroups that mainly carry binary data. It

has been estimated that, in terms of bytes, the top four newsgroups consume 22% of the entire volume. The top 35 groups consume 50% of the entire volume.

Currently, almost the only description of a message is its subject. This way of describing information items is more or less adequate for textual messages that contain text discussing the subject. For multimedia items, one-line descriptions can hardly be adequate. Normally, subject contains name of the collection or short description of the multimedia item, name of file, number of the part and total number of parts (such as "Persian kitten cats123.jpg (1/1) 35567 bytes"). This format is often used, but many multimedia postings do not have even that. Often subject lines are quite meaningless, e.g. "My loved kittens". Figure 1 shows a list of Usenet messages in a binary newsgroup alt.binaries.pictures.animals. Figure 2 shows a so-called "index" of files posted in the group. Clearly, the way of representing of available contents as it is shown on the Figure 2, is much more user friendly and informative for users.

Multimedia items, such as video, image, sound files, are bulky. It may take a minute to download an image over a slow modem line. For sound and video files, it may take several minutes and even hours to download them. This is why it is important to give users better means of selection, ensuring that they can make more informed downloading decisions and avoid wasting time and money on downloading large objects that will be discarded right after first viewing or hearing.

"Indices" or "contact sheets" are often posted with large collections of images. It is a requirement of "netiquette" to post indices when posting a collection of images. Such indices give users a chance to preview pictures of the collection before downloading them and decide whether they want to spend their time and resources on downloading these images, and/or select images that they want to download. Often small picture previews (frame captures) are posted with video files.

However, working with such indices is very inconvenient. Selecting and downloading desired files requires performing the following steps:

1. Find and download the index (thumbnail preview).
2. View the index. Identify file names of images represented by thumbnails.
3. Memorise or write down file names of images to be downloaded.
4. Locate messages that carry these files in the list of messages.
5. Download the found messages, extract and save files.

The system that we are inventing makes it very easy to achieve the same result. Users just download the index, view it, point to thumbnails of images that they desire to download and double click on the thumbnails to download related images or video files. We are inventing a way of use of index images that provides an easy, intuitive and efficient "point and click" interface for selecting and downloading multimedia items posted on the Usenet.

There are systems that offer to end users similar options to what we are targeting. There are sites on the Web that load binary files from the Usenet to their proprietary databases, create preview items (thumbnails for images, for example), and then give access to the contents of these databases in WYSIWYG mode using Web browsers. This is not what we are targeting. We are targeting a system that can give users an easy and convenient way of previewing and selecting multimedia items using a newsreader and working with a standard news server.

There are some patent applications targeting the same problem. A US Patent Application No 09/722676 with title "A Method and System for Communication in the Usenet" has been filed in 2000 by CSIRO Australia. In one of its aspects, the application describes a method and system for helping Usenet users to make informed decisions regarding downloading of multimedia items from the Usenet. In particular, the described application invents a method that allows to post metadata items associated with multimedia items, present these items for user preview, then locate and download multimedia items selected by the user based on the presented metadata items.

When used with image collections, this system allows to post small image thumbnails associated with messages that carry the image files, present these thumbnails to users for preview, then when users select one or more of the images for downloading, locate the needed messages and download them.

However, the invention disclosed in CSIRO application has 3 disadvantages that the present invention aims to eliminate.

1. It requires insertion of association information (tags) to headers of messages being posted. When such tags are automatically inserted in messages subjects, they make subjects look somewhat aesthetically unpleasing.
2. As a consequence of the requirement to insert tags in the messages headers, metadata items can only be posted with a new collection, by the software that generates and posts the messages. It is not possible to add metadata descriptions to a collection that has been already posted without them.
3. Images thumbnails are posted in a format that can't be used by newsreaders that are not specifically designed to work with this format, even if they can show images.

The present invention eliminates the disadvantages listed above. It does not require insertion of any additional/special information in messages carrying multimedia items. It enables adding metadata descriptions to collections that have been posted already by same poster or different people. And finally, it posts images thumbnails as standard indices (or contact sheets) commonly posted by Usenet users. The system includes additional information that allows associate thumbnails with messages that carry the actual images or videos.

Unfortunately, not everyone observes requirements of the etiquette. Many collections of images are posted without indices. There are groups where indices are rare. Video files are also very rarely posted with previews. Sometimes people help each other and previews are posted later, not by the person who posted the original articles, but by someone else who downloads them, builds previews and posts them for other people to use.

The system that we are inventing (the second aspect) makes it possible for Usenet users to request indices/previews for a particular set of articles. They send such requests to one or more servers that specialise in processing of such requests (index servers). These servers can be co-located with news servers, so that usage of Internet bandwidth required to build and post an index is minimal.

There are systems that offer to end users similar options to what we are targeting. There are sites on the Web that load binary files from the Usenet to their proprietary databases, create preview items (thumbnails for images, for example), and then give access to the contents of these databases in WYSIWYG mode using Web browsers. This is not what we are targeting. In the invented system, anyone can request an index for a particular set of articles, provided they have access to an index server. The index server builds this index, possibly caches it, and posts it to the Usenet or sends it back to the user who requested it. If the index has been posted to the Usenet, it becomes available for everyone for previewing and selecting multimedia items using a newsreader and working with a standard news server. If the service is configured so that it sends indices to users who requested them, then the users can use obtained indices to make better informed decisions on whether to download a particular object from the Usenet. Either way, this invention makes it possible for Usenet users to obtain information that helps them to choose what they want to download.

The third and fourth aspects of the present invention deal with advertising on the Usenet. The Usenet is being actively used for advertising. Many images and videos posted on the Usenet contain advertising references to Web sites. Many messages have URLs of advertised Web sites. However, unlike Web advertising, where users have just to point and click on a banner that attracts their attention, Usenet advertising is not as convenient to use. For example, if a URL is shown on an image posted on the Usenet, users have to re-type this URL in their browser to reach it. Needless to say that relatively few people go for so much trouble to visit an advertised site.

Enabling "point and click" functionality in Usenet advertising would certainly increase number of people who respond to advertising because pointing to a banner or URL and clicking on it is much easier and simpler than re-typing a URL in a browser. The other reason is that it allows designing and using for advertising graphical "clickable" objects specifically targeted at catching users attention and encouraging them to click on the object.

Another problem that the Usenet advertisement currently has is difficulty with counting referred visitors. Counting referred visitors is an important requirement in commercial advertising because advertising fees are often based on numbers of visits to the advertised Web site generated by the advertising. Counting such visits is relatively easy on the Web. For example, if site A is displaying banners advertising site B, it is not hard for site A to count user clicks on the advertising banners and thus the number of visitors referred to site B.

The situation is quite different in the Usenet environment. When users download Usenet messages, they are not in any contact with the party that has posted these messages. When users follow advertising and visit the advertised sites, no feedback of this activity is coming to the poster of the advertisement.

The second aspect of this invention deals with advertisement on the Usenet. It discloses a system that makes it easy ("point and click") to follow advertisement links. It also discloses a system that provides for accurate counting of visits to the advertised site generated by an advertisement posted on the Usenet.

3 Summary of Invention

3.1 First Aspect

A first aspect of the present invention provides a method of coordinating the identification of objects with their descriptions (metadata) in a newsgroup of the Usenet, the method including the steps of:

generating a graphical index (contact sheet) of images or videos posted in Usenet messages,
generating an image map of the graphical index that associates regions of the graphical index with Usenet messages,
posting the message containing the graphical index and the image map.

There is also provided a method of downloading messages from the Usenet, the method including the steps of:

receiving headers or only XOVER information of messages available for downloading,
scanning this received information to identify which messages contain graphical indices,
downloading the messages containing graphical indices,
presenting the indices to the user to make a decision regarding the downloading of associated objects,
if the user wants to download an associated object,
locating the image map record that corresponds to the region of the index where the user mouse is pointing to,
finding and downloading the message that the map record refers to.

It has been realised that images represent a significant part of multimedia objects posted on the Usenet. Users posting large collections of images (tens or hundreds of them) often post indices - images that contain thumbnails (small copies) of images posted in the collection or small copies of frames captured from video movies. This gives to the downloaders the opportunity to download first the index image and get a better idea about the images posted in the collection, make better informed decisions whether to download a particular image and thus save downloading time and money spent on the Internet session.

This aspect of invention is based on an automatic way of providing a graphical index for every multimedia item posted in a collection and associating thumbnails in the index with messages that carry the items these thumbnails represent.

This illustrates an approach that uses a different kind of description of an information item. Naturally, a little copy of image is a better description of it than a subject line.

This allows users to save downloading time and makes easier and more efficient the process of selecting and downloading media items. To download a set of selected images from a posted collection, currently, the user may have to perform the following steps:

locate collection messages;
locate collection indices;
download collection indices;
view collection indices and memorize (or write down) names of wanted files;
locate messages carrying the wanted files and either mark them for background downloading or download them instantly;

The present invention allows for establishing of connections between visible thumbnails in index images and messages containing the media objects that these thumbnails represent. The invention then uses these connections to identify and download media objects selected by the user. The user indicates their preferences preferably by pointing with their mouse pointer and clicking on thumbnails of objects that they want, when viewing the index image.

The connection is established in two stages: first, the thumbnail region is associated with a record in a special image map contained in the index message. Second, the record in the image map is associated with the message that has the media object (or at least a part of it) that is represented by the thumbnail.

The method of the first aspect preferably includes two stages.

Stage 1

A poster posts a collection of multimedia items. An index of the collection may be posted together with the collection or added later when the collection has been posted already.

A collection index is a message or several messages that include images containing thumbnails representing the collection items, and image maps for each of the index images. An image map is a set of records, each of which contains some information about a particular region of the image. Thus, if a thumbnail occupies a region, the record may include information that relates to the thumbnail and the item represented by the thumbnail. Records are associated with image regions by including in each record a set of coordinates that define a particular image region, as it is illustrated in Figure 3.

To find a record corresponding to a particular point in the image, records in the image map are searched and those that have coordinates of regions that contain the given point are identified.

Based on these records, the point then can be associated with a particular message or set of messages.

There are different ways to indicate an association between records in image map and corresponding messages carrying the collection items.

One of possible methods is described in the CSIRO patent application. It uses unique pairs of tags. One tag of the pair is included in the index message. Another tag of the pair is included in the header of the associated message. When the first tag is given, it is possible to construct the second tag and then find a message that has it in its headers. We described above the disadvantages that this method has.

Therefore, we use two other methods to achieve the same result (establish references from the index messages to the messages that carry media objects described by the index), but in a better way. Generally, any combination of message header fields (preferably fields included in XOVER) can be used, as long as it uniquely identifies the messages.

The first preferred method is used in a case when index with a new collection is posted. In this case no message-IDs of collection messages are available at the moment of posting. In this case, a pair <message subject, poster ID> is included in the image map record to identify the message the record refers to. It must be noted that some news servers return at the moment of posting message-IDs assigned to messages. Some other servers accept message-IDs generated by posting clients. Although it is possible to use these features, they can't be relied upon because they are not a part of NNTP.

The second preferred method is used in a case when index is posted for a collection that has been posted already and is available in one or more Usenet newsgroups. In this case unique message-IDs are available for all collection messages and can be included in image map records to establish references.

Stage 2

The downloader's client downloads headers of all new messages in the newsgroup or just XOVER information, which is sufficient if we use the preferred methods described above. The client may identify the index messages and automatically download them and present to the user, or it may do it in response to some action of the user,

indicating that the user wants to see the index. For example, the user may have to click on the index message in the message list in order to see the index.

At each point of time when the user moves their mouse pointer over the index image, the client receives mouse movement events and can track the pointer and is aware of the image region the pointer is in. When the user wishes to download/see a particular image/movie, they point to its thumbnail and do some action, e.g. double click or press enter key. The client determines the coordinates of the mouse pointer within the image and searches the image map for a region containing a point with these coordinates.

When such a record is found, the client reads either < subject, poster ID> pair or message-ID from this record, looks for a message that has these attributes, downloads it and presents to the user or does any other action that user wants to be done with this message.

The found message may contain only a part of the multimedia object wanted by the user, or be just an overview message. In such a case, messages that contain other parts of the object have similar subject that normally differ only by part numbers and are easy to locate. Modern binary newsreaders usually can do it automatically.

The user may have a number of options to select from when dealing with a particular thumbnail. Downloading is not the only possibility. For example, they can insert the media item into a list for automatic background downloading, delete corresponding messages from the message list, view the media item if it is already downloaded, etc.

This approach has been found to significantly simplify for users the process of selecting and downloading of multimedia items. For example, in case of images, user sees a set of small images (thumbnails) presented by the newsreader. Each of these images represents a "real" large image. To download real images, the user just has to double click on a thumbnail or select a few of them and then start batch downloading.

Although this kind of 'click on link' interaction is used on the Web, but the underlying protocol (HTTP) is different. Our invention makes it possible to achieve the same level of convenience when working with the Usenet-like systems.

The advantages of the present invention include:

- 1) A better representation of available messages during selection stage. This avoids downloading multimedia objects that are unwanted and will be discarded later anyway.
- 2) An index that is viewable in other newsreaders. Most newsreaders provide a way to view images attached to messages. Even if a newsreader is not designed to use image maps and locate referred messages automatically, the user still can use the index image to select and download wanted messages, although not as efficiently as this invention makes it possible.
- 3) Indices can be posted not only with a new collection, but also for existing collections, that have already been posted, possibly by other people.
- 4) This method has no effect on subject text or other fields of message headers. No ugly tags appear in headers.

This invention provides a general, flexible and easily extensible way of associating of additional information with messages and using this information when required.

3.2 Second Aspect

A second aspect of the present invention provides a method and system for providing "indices on request" service to the Usenet users.

The system includes a Usenet news server, a user client application (such as a newsreader) and an index server, all three components able to communicate with each other via network connections.

The method includes the steps of:

the user identifying a set of articles that they would like to be indexed;

the user submitting information identifying these articles to the index server;

after receiving and accepting this request, the index server downloading these articles, constructing their previews (indices) if possible and either submitting these indices to the news server, or sending them back to the user client application, or both;

optionally, the user downloading and viewing the constructed indices.

This gives to the downloaders the opportunity to download first the preview (index) information and get a better idea about the multimedia items available for downloading, make better informed decisions whether to download a particular item and thus save downloading time and money spent on the Internet session.

This invention minimises users efforts and expenses involved in publishing indices. All users have to do is to compose a list of articles carrying the multimedia items to be indexed and submit this list to an indexing server that takes care of retrieving the items, constructing their previews and publishing (posting) these previews on the news server. Because these two servers can be co-located, all the traffic involved in retrieval of bulky multimedia items from the news server in order to construct their previews, would be local. This insures that the system is efficient and economic.

Once indices (previews) have been published and are available, they significantly decrease amounts of wasted time and resources spent on downloading of unwanted material.

It is technically possible to implement completely automatic indexing of multimedia newsgroups, but the quality of service would be relatively low, for at least two following reasons:

1. It is impossible to determine automatically whether a collection of items to be indexed is appropriate (on topic) for a particular newsgroup. Indexing inappropriate collections and advertising junk posts would constitute violation of Usenet etiquette because indices/previews of off-topic items would be also off-topic.
2. In many cases, it is hard to determine automatically whether a collection already has an index included. Creating additional indices would lead to unjustified spending of Usenet resources.

This invention also discloses a method for automatically identifying collections of images.

3.3 Third Aspect

A third aspect of the present invention provides a method of using images for advertisement on the Usenet, the method including the steps of:

generating an image containing advertisement regions,
 generating an image map of the image that associates regions of the image with URLs of Web pages,
 posting the message containing the image and the image map,
 downloading the messages containing advertising images,
 presenting the images to the user to view,
 if the user wants to follow advertisement,
 locating the image map record that corresponds to the region of the index where the user mouse is pointing to,
 accessing the URL, preferably using a Web browser, if a URL found in the image map record.

The present invention allows for establishing of connection between image regions and URLs they are advertising. The invention then uses these connections to access URLs associated with image regions selected by the user. The user indicates their preferences preferably by pointing with their mouse pointer and clicking on advertising parts of the image.

The connection between the advertising part of the image and the target Web site is established in two stages: first, the image region is associated with a record in the image map contained in the index message. Second, the record in the image map is associated with the Web page by including its URL.

The method of the second aspect preferably includes two stages.

Stage 1

Poster posts an advertising message. An advertising message contains an image containing graphical advertising and an image map of this image. An image map is a set of records, each of which contains some information about a particular region of the image. Thus, if a region is advertising a particular Web site, the record may include the URL of the site. Records are associated with image regions by including in each record a set of coordinates that define a particular image region.

To find a record corresponding to a particular point in the image, image map records are searched and those that have coordinates of regions that contain the given point are identified.

Based on these records, the point (or region) then can be associated with a particular URL.

Stage 2

The user client downloads an advertising message and presents the image contained in the message to the user.

At each point of time when the user moves their mouse pointer over the index image, the client receives mouse movement events and can track the pointer and is aware of the image region the pointer is in. When the user wishes to follow a particular advertisement, they point to the region of the image that represents it and indicate their wish, e.g. double click or press a key. The client determines the coordinates of the mouse pointer within the image and searches the image map for a region containing a point with these coordinates.

When such a record is found, the client initiates a sequence of actions leading to visiting the URL that is being advertised.

The advantages of the present invention include the ease of accessing advertised sites. This will result in significantly increased response to advertising efforts, i.e. number of visitors to advertised sites generated by advertisement posted on the Usenet.

3.2 Fourth Aspect

A fourth aspect of the present invention provides a method of counting visits to sites advertised on the Usenet, the method including the following steps:

establishing a Web site or a script on a Web site,
generating a URL that has a direct or indirect reference to the advertised site,
posting the generated URL in advertisement messages,
when a user accesses the URL, processing its parameters, increasing visitors count and referring the user to the advertised site.

This method is based on use of a dedicated Web site or a script on a Web site to count visitors and refer users to advertised sites. We will call this site/script a "referring site" or "redirecting site".

When a user accesses such an advertisement URL, their browser first comes to the referring site and passes to the referring script the parameters that identify the target (advertised) site. The script identifies the target site, increases the counter of visitors of this site and refers the browser to the target site. Referring a Web browser from one URL to another is not a problem and is routinely done on the Web.

Examples or advertising URLs: <http://sendmeto.com/goto?url=http://www.mysite.com> or
<http://sendmeto.com/goto?account=#42527829090>

In both of these examples, "sendmeto.com" is a referring site and "goto" is name of the script on this site that does counting and refers visitors to the target site.

In the first example, the target URL is included in the advertising URL explicitly. In the second example, the target URL can be identified via the account (profile) number included in the advertising URL. The processing script has to find the account and get the target URL from there in order to refer the visitor.

The present invention provides a way to establish accurate counting of visits generated by advertisements posted on the Usenet.

4 Embodiment

4.1 First Aspect

A practical implementation of this invention does not require changing of involved standards, such as NNTP, MIME etc. It only requires modification of posting and downloading news clients so that they would add some extra information to messages carrying collection indices and could interpret this information during the downloading stage.

To describe how the system works, we will take as a base work of Ozum newsreader, version 2.86 build 1107 that has in-built functionality to automatically produce standard contact sheets (thumbnail sheets, indices) when posting collection of images. We will refer to this newsreader as "Ozum" in the text below. Those skilled in the art are familiar with use of a typical news client. We will describe how the improved client works in our embodiment. To do this, we will describe what it does differently or additionally to Ozum.

For sake of simplifying this description, we will call OzumPlus an improved Ozum newsreader with functionality disclosed in this invention.

The format of index messages generated by OzumPlus has to be such that when a user clicks or double clicks on any thumbnail within the index image, the information included in the index message would allow identifying the related message (for a thumbnail). This can be achieved by including an image map in the index message, as Figure 3 illustrates it.

When posting indices with new collections or posting indices for collections that have been posted already, OzumPlus has to generate indices as Ozum does it now, plus include this additional information.

4.1.1 Format of Index Messages

Index messages include index images and their image maps (preferably placed after the image). An image map is preferably a set of lines defining associations between image regions (preferably rectangles) and messages that carry media items represented by the thumbnails or URLs corresponding to the banners. Image map preferably starts with a line containing a string: `<image_map>\n` and ends with a line containing a string `</image_map>`. Any other distinctive strings can be used as long as they are unlikely to occur in messages that don't have image maps.

For every thumbnail, the image map includes a line in format similar to standard (Web, HTML) image map format:

```
rect <<message identification info> crc=<crc> size=<size>> coordinates
```

where

rect - is a method name, serving for conceptual compatibility with image maps and for possible later extensions.

message identification info - is either message id of the message referred by the map line, or `<subject, poster ID>` pair. It is best to use message ids for purpose of identifying messages, but they are not available, for example, when we are posting a new collection and want to post indices with it. In this case `<subject, poster ID>` pair has to be used and file name or number included in the subject in order to ensure uniqueness of subjects within the collection.

coordinates – preferably left top and right bottom coordinates of the rectangle in the index image occupied by the thumbnail. Clicking (or double clicking) in this rectangular area of the image will cause downloading of the associated message. Although rectangular areas are used in this embodiment, other shapes can be used too.

crc and size are optional fields for the purpose of this invention, and can be used to determine that an identical object has been downloaded by this newsreader in the past, or for any other purposes, such as verification of correctness of downloading and decoding.

If index image is large and is split in several parts when posting, a full image map is preferably posted with every part.

4.1.2 Algorithms

4.1.2.1 Posting Indices with a New Collection

Creating indices happens when people send indices with new collections. Message ids are not available at the time of posting, so, other information, less reliable than message ids, must be used.

OzumPlus creates and posts index as Ozum does it, except:

For every thumbnail put an entry in the image map containing the thumbnail rectangle coordinates and <subject, poster> pair preferably in format:

rect <subject=subject text ‘t’ poster=poster text crc=<crc> size=<size>> coordinates

Here ‘t’ is a tab character. It can not occur in subject or poster text and therefore is convenient for use as a separator.

4.1.2.2 Adding Indices to an Existing Collection

Adding indices happens when a user wants to add an index to a collection that has been posted already. In this case they select collection messages in the messages list, right-click on the selection and choose “Add Index” in the context menu.

OzumPlus downloads the messages, extracts media items, creates and posts the indices.

This case is similar to the case of posting indices with a new collection, except in case of adding indices, as opposed to the case of creating indices, message ids of indexed messages are available and can be used instead of their <subject, poster> pairs. Message ids are more reliable for messages identification and therefore preferable.

So, the entries in image map are preferably in following format:

rect <msgid=<msg-id> crc=<crc> size=<size>> coordinates

4.1.2.4 Using Indices

When user double clicks on a message containing an image, OzumPlus checks whether this message also contains an image map.

If the image and/or image map are split into several parts and posted in several messages, OzumPlus downloads all the messages and assembles the image and/or the image map. This procedure is done routinely now by many newsreaders and does not represent a significant technical problem.

OzumPlus reads the image map and verifies that it reads OK and all its parameters such as coordinates of rectangles are in reasonable ranges. If not, the image map is ignored.

If the image map is recognized, the regions that it defines in the image are treated as “clickable”. OzumPlus reads the map and preferably stores its entries in a list in the memory before displaying the image.

When user moves their mouse pointer over the index image, OzumPlus receives and preferably processes mouse movement events in order to identify the image map entry that corresponds to the point over which the mouse pointer is. To do this, OzumPlus uses coordinates of the areas included in the image map entries.

When the user puts the mouse pointer over some point in the image and does something, for example, presses a particular key or mouse button, OzumPlus does the following

1. If there is no image map entry identified for the point and the action requires such entry, do nothing.
2. If there is an entry found for this point, identify the message based on the information available in the entry and perform the function corresponding to the action done by the user. For example, if the user double clicks, download the message and present it to the user.

A preferable scenario of using index images and image maps is described below.

OzumPlus displays the index image and traces mouse movement over the image. When a click or double-click event comes and the mouse pointer is over a rectangle defined in the image map:

1. If the rectangle corresponds to a thumbnail and the event is a single left mouse button click, use message-ID or <subject, poster> pair that provided in the image map to identify the message and:
 - a. If the message is marked for downloading – unmark it.
 - b. If the message is not marked – mark it.
 - c. If the message is downloaded or not available – do nothing.
 - d. If the message is not marked and, based on crc and size, it is clear that this message has been downloaded before – display a dialog asking for downloading confirmation.
2. If the rectangle corresponds to a thumbnail and the event is a double left mouse button click – use message-ID or <subject, poster> pair that provided in the image map to identify the message and download and display the message.
3. If the event is a double right click – do nothing.
4. If the event is a single right click – display a menu with the following options:
 - a. Mark to Download (disabled if the message is marked already);
 - b. Unmark (disabled if the message is not marked);
 - c. Download Now (disabled if the message is downloaded or not available);
 - d. Mark as Read;
 - e. Select (add to a list of currently selected messages in the group list).
5. After the user chooses any of the actions available from the menu use message-ID or <subject, poster> pair that provided in the image map to identify the message, then perform on it the action requested by the user.

4.2 Second Aspect

A practical implementation of this invention involves implementing the client side and the server side.

Functions performed by the client side:

1. Allowing the user to provide address and authentication information necessary to connect and interact with the index server.
2. Allowing the user to browse articles available for downloading from the news server.
3. Allowing the user to select a number of articles and indicate that an index (preview) is to be built for items carried by these articles.
4. Allowing the user to input parameters of indexing. For image indices such parameters are numbers of columns and rows of thumbnails on page, size of thumbnails, etc.
5. Forming text of the index request.
6. Connecting to the index server and sending the index request there.

7. Presenting the server's response to the user.

To describe how the system works, we will take as a base work of Ozum newsreader, version 2.87 build 1111 that has in-built functionality to produce standard contact sheets (thumbnail sheets, indices) for images available for downloading. We will refer to this newsreader as "Ozum" in the text below.

Those skilled in the art are familiar with use of a typical news client. We will describe how the improved client works in our embodiment. To do this, we will describe what it does differently or additionally to Ozum. For sake of simplifying this description, we will call OzumPlus an improved Ozum newsreader with the user client functionality disclosed in this invention.

Using Ozum, it is possible to select a number of articles and order Ozum to download these articles, build index of images carried by these articles, and post this index back to the news server. Thus, Ozum provides an easy way for adding missing previews to collections of images. The disadvantage of this method is that the user who orders indexing cannot benefit from it because building of indices requires downloading of all articles that are being indexed. So, by the time the index is ready and published on the news server, all the indexed items are downloaded and stored on the computer of the user who ordered indexing.

Other users, however, can use the newly added index to choose articles that they want to download. Ozum can be modified to not to create indices locally, but instead form and submit an indexing request to an index server.

To implement the required functionality, the following modifications can be implemented:

1. An additional "Index Server Properties" dialog allowing users to input index server address, port number and authentication information (user name and password).
2. The functions that create index have to be modified so that after indexing job has been formed ("Create Index" menu option in Ozum), OzumPlus submits the list of articles and job parameters to the index server in a certain format discussed below. After submission, OzumPlus receives response from the server and reports to the user whether the index order has been accepted or there is a problem.

Both modifications are not hard to implement for those skilled in the art. An extended subset of NNTP commands can be used for communications between OzumPlus and index servers. This set includes standard NNTP commands AUTHINFO USER, AUTHINFO PATH and QUIT. An additional XINDEX command is introduced that, when issued by the client, asks the server to accept an index order information. The server either agrees to receive it or rejects. This works similar to the NNTP POST command.

Example 1.

Client: XINDEX
Server: 340 OK, send your data.
Client: <... Request data followed by a line containing a single dot.>
Server: 240 request accepted.

Example 2.

Client: XINDEX
Server: 440 Sorry, too busy, accepting new requests is temporarily suspended.

Example 3.

Client: XINDEX
Server: 340 OK, send your data.
Client: <... Request data followed by a line containing a single dot.>
Server: 441 request rejected, bad request data.

Index request data contains a list of articles to index and all other information that is necessary to build indices, such as attributes if fonts used, size of thumbnails, etc. XML-like encoding can be used to encode request data and send it to index server. Such encoding allows easily extracting data items from the request text, using available XML software to do that, and also adding new types of items to requests if necessary.

Example of encoded request data:

```

<header text="This is index header text." font="Arial" size="18" style="bold">
<footer text="This is index footer text." font="Georgia" size="14" style="bolditalic">
<captions font="Arial Narrow" size="10">
<thumbnails size="100" rows="3" columns="6">
<newsgroup name="alt.binaries.multimedia">
<message-ids>
hdsgghsk1653j@news.company.com      12325426
hdsgghsk16534j@news.company.com      12325427
hdsgghsk16535j@news.company.com      12325428
hdsgghsk16538j@news.company.com      12325429
...
</message-ids>

```

Articles numbers in the group on the server are preferably included in addition to textual message ids, for reliability of retrieval. Theoretically, articles numbers are not needed because news servers are supposed to be able to find and serve articles based on textual ids only. In practice, this does not always work. A server may respond that the article is unavailable when given a textual id, and yet successfully find it when given the number of the article in the group.

Example of a complete session.

Client connects to server.

Server: 200 Hello, this is index server at mynews.com.

Client: AUTHINFO USER joke

Server: 381 user name OK, send password

Client: AUTHINFO PASS notfunny

Server: 281 authentication accepted

Client: XINDEX

Server: 340 OK, send your data.

```

Client: <header text="This is index header text." font="Arial" size="18" style="bold">
<footer text="This is index footer text." font="Georgia" size="14" style="bolditalic">
<captions font="Arial Narrow" size="10">
<thumbnails size="10" rows="3" columns="6">
<newsgroup name="alt.binaries.multimedia">
<message-ids>
hdsgghsk1653j@news.company.com      12325426
hdsgghsk16534j@news.company.com      12325427
hdsgghsk16535j@news.company.com      12325428
hdsgghsk16538j@news.company.com      12325429
...
</message-ids>

```

Server: 240 request accepted.

Client: QUIT

Server and client close the connection.

Functions performed on the server side:

1. Listening for incoming connection requests.
2. Accepting connections, authenticating clients.
3. Accepting or rejecting index requests.
4. Performing index requests.

The first three functions are trivial. Almost any server performs similar functions. The only difference is in request types. After receiving an index request, the server should check its format and check that values of all parameters are valid.

The system consisting of the client and server can be modified to work so that the created by the index server index is returned directly to the client, instead of, or in addition to, being posted to the Usenet. In this modification, the index server preferably maintains a cache of indices created by it, to optimise the performance.

In case if produced indices are being posted to the Usenet, in addition to, or instead of, returning them to the requesting user, to avoid wasting of Usenet resources and ensure a higher quality of service, index servers should not repetitively post indices of articles that have been indexed already. For this purpose, index server can maintain a list of articles that have been indexed in the last few days. The server does not post indices of articles that are in the list of indexed or are too old (e.g. older than the oldest article in the list).

It is harder to avoid repetitive indexing when there are many index servers run in parallel, possibly by different companies. In this case, one of them (let us call it "Server 1") can maintain a global list of indexed articles. Other servers, before posting an index, can connect to Server 1, send it a copy of the request (or just a list of articles to be indexed) in order to verify that articles that they are asked to index, have not been indexed yet by any of index servers. After receiving a verification request from another server, Server 1 checks the global list of indexed articles and replies whether the index should be posted. If the articles have not been indexed yet, it recommends posting of indices and adds the articles to the global list of indexed articles.

It is also possible to avoid repetitive indexing and posting by introducing specialisation among index servers. Every server is assigned a subset of Usenet newsgroups that it can index. It accepts only index requests concerning these newsgroups and rejects index requests involving indexing material from newsgroups outside of the assigned subset.

After the index request has been verified and accepted, the index server that accepted it either finds ready indices in the cache, or downloads articles that it was asked to index, constructs indices or previews of attached multimedia objects and optionally posts these previews to a newsgroup where users can find them. It can be the original newsgroup(s) the articles were downloaded from, or it can be a special newsgroup dedicated to posting indices/previews of multimedia objects.

It must be noted that index server and news server don't have to be separate instances. Index server functions can be built in into a news server. This system would be even more efficient than separate index server and news server.

As it was noted above, newsgroups can be indexed completely automatically. That is, indices/previews of multimedia items are created soon after the items become available in the group.

Constructing samples of video and audio material is relatively easy. Articles carrying all or just first few parts of the file must be downloaded, the file assembled and then samples/previews constructed by an applicable for this media type method. For example, for a video file, some frames can be extracted and converted to images, or a much shorter video file can be constructed from short video sequences taken from the original file.

It is harder to achieve acceptable quality of indexing for images. There is no sense in indexing of individual images. To achieve an acceptable quality of service, only collections of images should be indexed. We invent a heuristic method for identifying posted collections of images for the purpose of indexing. A set of articles is determined to carry a collection of images if it satisfies to the following conditions:

1. The articles have images attached. This is very likely if articles' subjects include file names with extensions of image types, such as .gif or .jpg.
2. They were posted by the same poster. A number of criteria can be used to establish this. First, contents of the "From" field of the articles must be identical. Second, contents of other (some of them optional) header fields, such as "Path", "X-Trace", "NNTP-Posting-Host", etc should be identical or at least similar.
3. Articles must have identical comment part of the subject. That is, the part of the subject that remains after removing file names, size information, file number in the post, CRC, etc.
4. Articles were posted within a relatively short interval from each other. If article B follows article A, then time of posting B is not later than time of posting A plus some preset interval. For example, it is very rare that someone would post a part of collection and then continued posting the same collection after an hour interval. Most collections are posted in bulk, articles following each other within minutes or even seconds.

After a set of articles has been determined to be a collection using the criteria described above, the articles can be automatically downloaded, images extracted, indices built and posted.

4.3 Third Aspect

The system disclosed for the first aspect, has to be only slightly modified to enable posting and processing of advertisements.

1. It must allow construction of images where some regions are advertising regions, such as banners. This does not represent a technical problem and can be done, for example, by inserting one or more advertisement banners in the images.
2. It must allow association of the advertising regions in the image with URLs of the advertised Web sites. Introducing a slightly different type of image map records does this. Records of this type contain a URL instead of message identification information, preferably in format:

rect <URL="url-text"> coordinates

3. Algorithm of processing of image map records has to be slightly expanded to accommodate for processing of advertisements. In particular, when the user clicks on an image and the newsreader finds that this region corresponds to a map record with a URL in it, the newsreader must know that the user wants to follow an advertisement link and react correspondently. Preferably, the newsreader will start a Web browser and direct it to the URL specified in the map record.

4.4 Fourth Aspect

For the purpose of counting visits, this invention requires a set up of a special Web site (or just of a script on an existing site) that redirects Web requests based on their parameters and at the same time maintains counters of visitors redirected to particular sites.

Suppose, the address of such a site is <http://sendmeto.com/> and the name of the script that processes redirection requests is goto. The site has a database with containing URLs of target sites and associated visitor counters and possibly account (or profile) numbers.

This script does the following:

1. Parse received request and extract its parameters.
2. If the parameters contain a target URL:
 - a. Find in the database a record with this URL.
 - b. If such record does not exist, reply with an error message or just redirect the browser to the target URL.
 - c. If the record exists, increment the associated visitor counter, update the record and redirect the browser to the target URL.
3. If the parameters contain an account number:
 - a. Find in the database a record with this account number.
 - b. If such record does not exist, ignore the request or reply with an error message.
 - c. If the record exists, increment the associated visitor counter, update the record and redirect the request to the associated URL.

Having set up such redirection Web site and/or script, we can use for advertisement URLs that bring visitors to advertised sites not directly, but via the redirection site. We do this by generating URLs that pass the URLs of advertised sites as parameters to the processing script at the redirection site.

For example, if we want to send visitors to <http://mysite.com/>, in our advertisement messages we use URL

<http://sendmeto.com/goto?url=http://mysite.com/>

For example, we put `http://mysite.com/` in advertising message. When newsreader finds it, it automatically changes it to `http://sendmeto.com/goto?url=http://mysite.com/` and starts a Web browser with this URL.

6. Optionally, the user downloads articles chosen based on the indices.

6 The Claims Defining the Invention Are as Follows

1. A method of coordinating the identification of media objects with their associated indices in a newsgroup of the Usenet, the method including the steps of:
 - generating an index image (contact sheet) of the media objects,
 - generating an image map of the index, the image map associating thumbnails in the index with messages carrying media items corresponding to the index thumbnails,
 - generating and posting messages containing index images and their image maps.
2. A method as claimed in claim 1 where index map entries are associated with thumbnails in the index image via coordinates of the thumbnails rectangles.
3. A method as claimed in claim 1 where index map entries are associated with arbitrary image regions via coordinates of the image regions.
4. A method as claimed in claims 1-3 where image map entries are associated with newsgroup messages by including message-ids of the messages in the map entries.
5. A method as claimed in claims 1-3 where image map entries are associated with newsgroup messages by including any combination of messages header fields in the map entries.
6. A method as claimed in claims 1-5 where index image or any of the collection items may be split into several parts and each part posted in a separate message.
7. A method as claimed in claims 1-6 where index is generated and posted with a new collection.
8. A method as claimed in claims 1-6 where index is generated and posted for a collection that has been posted before and possibly by other poster.
9. A method of downloading messages from the Usenet, the method including the steps of:
 - receiving headers or only XOVER information of messages available for downloading,
 - downloading the messages containing indices,
 - presenting the indices to the user to make a decision regarding the downloading of associated objects, if the user wants to download an associated object,
 - finding in the image map a record corresponding to the thumbnail selected by the user,
 - finding and downloading the message that is associated with the image map record.
10. A method as claimed in claim 9 where index map entries are associated with thumbnails in the index image via coordinates of the thumbnails rectangles.
11. A method as claimed in claim 9 where index map entries are associated with arbitrary image regions via coordinates of the image regions.
12. A method as claimed in claims 9-11 where image map entries are associated with newsgroup messages by including message-ids of the messages in the map entries.
13. A method as claimed in claims 9-11 where image map entries are associated with newsgroup messages by including any combination of messages header fields in the map entries.
14. A method as claimed in claims 9-13 where index image or any of the collection items may be split into several parts and each part posted in a separate message.
15. A method as claimed in claims 9-14 where index was generated and posted with a new collection.
16. A method as claimed in claims 9-14 where index was generated and posted for a collection that has been posted before and possibly by other poster.

17. A method of associating of a region of an image posted on the Usenet with a Usenet message or a Web site, the method including the steps of:
- generating an image map of the image,
 - inserting in the image map a record containing coordinates of the image region to be associated,
 - inserting in the same image map record an identifier of the object that the image region is to be associated with,
 - posting the image and the image map in a Usenet message.
18. A method as claimed in claim 17, where a message-ID is used to identify a Usenet message.
19. A method as claimed in claim 17, where a <subject, poster ID> pair is used to identify a Usenet message.
20. A method as claimed in claim 17, where any combination of message header fields or their parts is used to identify a Usenet message.
21. A method as claimed in claim 17, where a URL is used to identify a Web site.
22. A method as claimed in claim 17, where any identifier is used to identify an addressable object.
23. A method as claimed in claim 17, where the image and/or the image map are split in parts and posted in several messages.
24. Using a method as claimed in claim 17 for creating and posting messages containing advertisement images, regions of which are associated with URLs of advertised Web sites.
25. A method of processing messages containing advertisement images, regions of which are associated with URLs of advertised sites, the method including the steps of:
- downloading and presenting the message and the image to the user,
 - if the user points to any point of the image and clicks or double clicks on it, or expresses a request for action in any other established way, finding in the image map a record corresponding to the point,
 - if the record is found and it contains URL of a Web site starting a Web browser and directing it to the target Web site.
26. A method as claimed in claim 25 where the Web browser is directed not straight to the target Web site, but first to a redirection site that will direct it to the target Web site.
27. A method of counting visits generated by advertisement messages posted on the Usenet, the method including the steps of:
- setting up a redirecting Web site or only a script for visits counting,
 - generating a URL pointing to the redirecting site and including information allowing to locate the advertised target site,
 - when the user wants to access the target site, directing their browser first to the redirecting Web site,
 - incrementing visitor counter maintained at the redirecting site,
 - redirecting the user browser to the target Web site.
28. A method as claimed in claim 27 where the information allowing to locate the advertised target site is its URL.
29. A method as claimed in claim 27 where the information allowing to locate the advertised target site is a key of a database record (e.g. account number) that has, or allows to establish, URL of the target Web site.
30. A method as claimed in claims 27-29 where the URL pointing to the redirection site is generated and inserted into the advertising message by the poster or posting software.
31. A method as claimed in claims 27-29 where the user news client automatically generates the URL pointing to the redirection site based on the original URL found in the advertisement message.

32. A method and system for providing on request previews (metadata descriptions, indices) of multimedia objects posted on the Usenet, the system consisting of:
- A functioning news server;
 - A user NNTP client application (newsreader);
 - An index server;
- The method including the steps of:
- Using the news reader, choosing a set of articles to be indexed;
 - Forming index request and submitting it to the index server;
 - Downloading the articles by the index server;
 - Constructing indices/previews by the index server and posting them to the news server.
33. A method and system for providing on request previews (metadata descriptions, indices) of multimedia objects posted on the Usenet, the system consisting of:
- A functioning news server;
 - A user NNTP client application (newsreader);
 - An index server;
- The method including the steps of:
- Using the news reader, choosing a set of articles to be indexed;
 - Forming index request and submitting it to the index server;
 - Downloading the articles by the index server;
 - Constructing indices/previews by the index server and either posting them to the news server or sending them back to the user who requested them, or both, depending on implementation and configuration.
34. A system and method as in claim 32 where index server functions are built in into the news server.
35. A system and method as in claim 33 where index server functions are built in into the news server.
36. A system and method as in claims 32 and 34 where the index server maintains a list of articles indexed before, in order to avoid repetitive indexing.
37. A system and method as in claims 33 and 35 where the index server maintains a list of articles indexed before, in order to avoid repetitive indexing.
38. A system and method as in claims 32 and 34 where the system consists of multiple news servers and index servers.
39. A system and method as in claims 33 and 35 where the system consists of multiple news servers and index servers.
40. A system and method as in claims 32, 34 and 38 where one or more of the index servers maintain global lists of indexed articles in order to avoid repetitive indexing. Other servers verify requests with them before processing. The servers maintaining the lists exchange updates as soon as new articles are added to the lists.
36. A system and method as in claims 33, 35 and 39 where one or more of the index servers maintain global lists of indexed articles in order to avoid repetitive indexing. Other servers verify requests with them before processing. The servers maintaining the lists exchange updates as soon as new articles are added to the lists.
37. A system and method as in claims 32, 33 and 35 where each index server is specialised and dedicated to creating indices only for a particular subset of newsgroups.
38. A system and method as in claims 33, 35 and 39 where each index server is specialised and dedicated to creating indices only for a particular subset of newsgroups.
39. A method for identifying posted image collections, in order to index them automatically. The method determines articles to carry an image collection if they carry images, have been posted by the same poster, share subject comment part and have been posted within a reasonably short time interval from each other.

Drawings

Date	Size	Subject	From
12/09/02 ...	71,661	Animals - 008182b.jpg(1/1) 51018 bytes	buffolo@mail.com
12/09/02 ...	89,632	Animals - 014009b.jpg(1/1) 63854 bytes	buffolo@mail.com
12/09/02 ...	60,561	Animals - 014010b.jpg(1/1) 43088 bytes	buffolo@mail.com
12/09/02 ...	60,096	Animals - 014012b.jpg(1/1) 42754 bytes	buffolo@mail.com
12/09/02 ...	84,056	Animals - 014039b.jpg(1/1) 59869 bytes	buffolo@mail.com
12/09/02 ...	37,188	Animals - 014043b.jpg(1/1) 26392 bytes	buffolo@mail.com
12/09/02 ...	91,640	Animals - 016055b.jpg(1/1) 65289 bytes	buffolo@mail.com
12/09/02 ...	103,149	Animals - 038092b.jpg(1/1) 73509 bytes	buffolo@mail.com
12/09/02 ...	74,523	Animals - 041044b.jpg(1/1) 53061 bytes	buffolo@mail.com
12/09/02 ...	80,701	Animals - 041048b.jpg(1/1) 57474 bytes	buffolo@mail.com
12/09/02 ...	75,425	Animals - 041053b.jpg(1/1) 53706 bytes	buffolo@mail.com
12/09/02 ...	85,840	Animals - 041062b.jpg(1/1) 61146 bytes	buffolo@mail.com
12/09/02 ...	98,853	Animals - 041067b.jpg(1/1) 70438 bytes	buffolo@mail.com

Figure 1. A part of a typical list of newsgroup messages




Figure 2. An index image

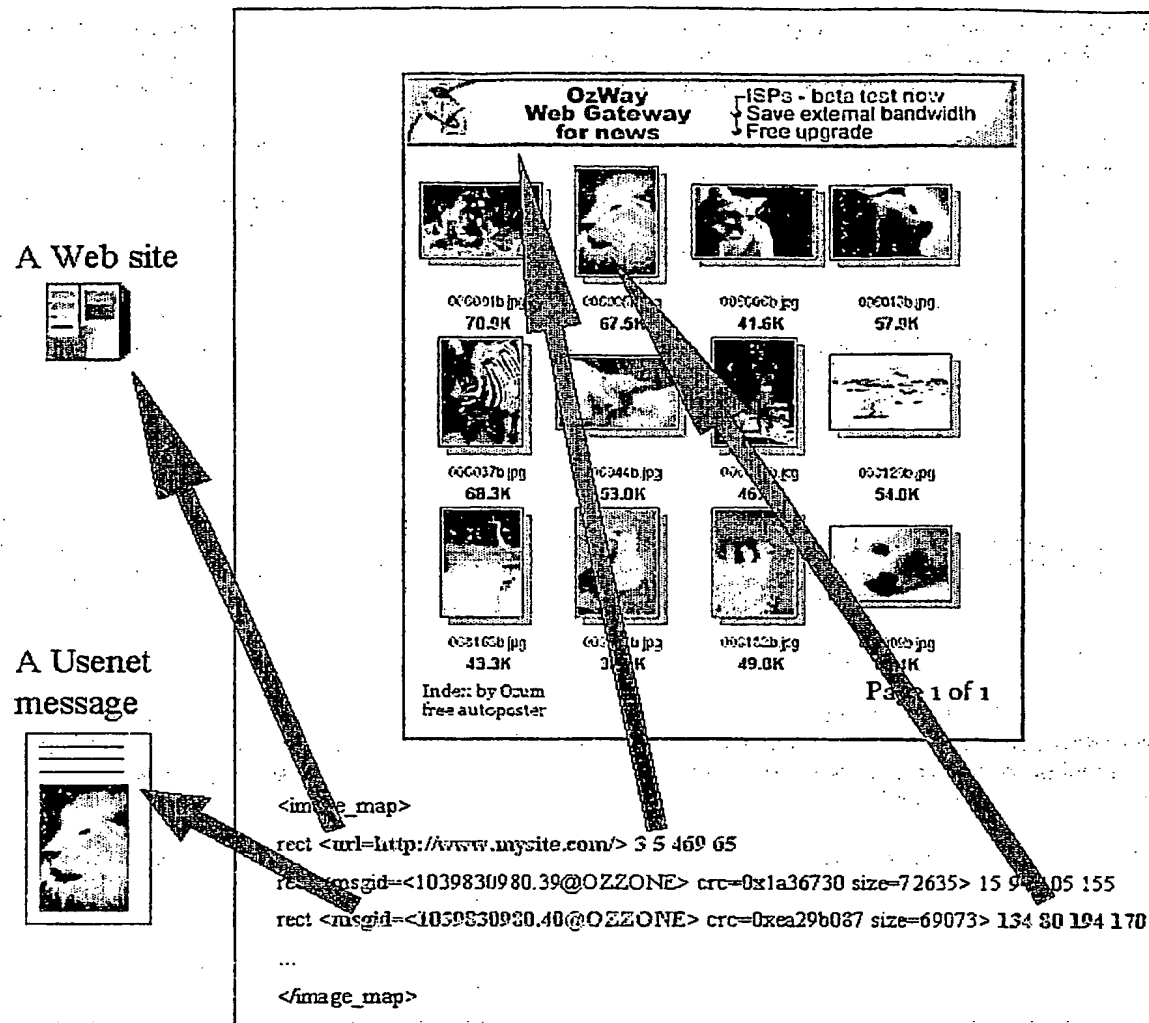


Figure 3. Image map records establish associations between image regions and Usenet messages or Web sites




Figure 4. Data flow in a system with one index server




Figure 5. Data flow in a system with multiple index servers

INTERNATIONAL SEARCH REPORT

International application No.

PCT/AU2003/001668

A. CLASSIFICATION OF SUBJECT MATTER				
Int. Cl. ⁷ : G06F 17/30				
According to International Patent Classification (IPC) or to both national classification and IPC				
B. FIELDS SEARCHED				
Minimum documentation searched (classification system followed by classification symbols)				
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched				
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) USPTO, DWPI (usenet, rntp, newsgroup, image, picture, binaries, thumbnail, image map, index, preview, summary, etc.)				
C. DOCUMENTS CONSIDERED TO BE RELEVANT				
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.		
X	WO 2001/050337 A1 (COMMONWEALTH SCIENTIFIC AND INDUSTRIAL RESEARCH ORGANISATION), 12 July 2001 page 17, lines 1-31; page 41, line 5 – page 49, line 11	1-24		
A	US 6,493,703 B1 (KNIGHT et al), 10 December 2002 the whole document	1-24, 32-43		
A	US 5,815,663 A (UOMINT), 29 September 1998 the whole document	1-24, 32-43		
<input type="checkbox"/> Further documents are listed in the continuation of Box C <input checked="" type="checkbox"/> See patent family annex				
<table style="width: 100%; border: none;"> <tr> <td style="width: 50%; vertical-align: top;"> <p>* Special categories of cited documents:</p> <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier application or patent but published on or after the international filing date</p> <p>"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p> </td> <td style="width: 50%; vertical-align: top;"> <p>"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention</p> <p>"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone</p> <p>"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art</p> <p>"&" document member of the same patent family</p> </td> </tr> </table>			<p>* Special categories of cited documents:</p> <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier application or patent but published on or after the international filing date</p> <p>"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p>	<p>"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention</p> <p>"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone</p> <p>"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art</p> <p>"&" document member of the same patent family</p>
<p>* Special categories of cited documents:</p> <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier application or patent but published on or after the international filing date</p> <p>"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p>	<p>"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention</p> <p>"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone</p> <p>"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art</p> <p>"&" document member of the same patent family</p>			
Date of the actual completion of the international search 23 January 2004		Date of mailing of the international search report - 5 FEB 2004		
Name and mailing address of the ISA/AU AUSTRALIAN PATENT OFFICE PO BOX 200, WODEN ACT 2606, AUSTRALIA E-mail address: pct@ipaustalia.gov.au Facsimile No. (02) 6285 3929		Authorized officer M. D. HOLLINGWORTH Telephone No : (02) 6283 2024		

INTERNATIONAL SEARCH REPORT

International application No.
PCT/AU2003/001668

Box I Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos :
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claims Nos :
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. ☐ Claims Nos :
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a)

Box II Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

See extra sheet.

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims
2. ☐ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:
1-24, 32-43

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest.
☐ No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

International application No.

PCT/AU2003/001668

Supplemental Box

(To be used when the space in any of Boxes I to VIII is not sufficient)

Continuation of Box No: II

The international application does not comply with the requirements of unity of invention because it does not relate to one invention or to a group of inventions so linked as to form a single general inventive concept. In coming to this conclusion the International Searching Authority has found that there are four inventions:

- Claims 1-24 and 32-43*, directed to methods and systems for indexing Usenet messages containing media objects. The automated generation of message indices is considered to constitute a first "special technical feature."
- Claims 25-26, directed to a method of processing a message containing an advertising image. The use of an image map to direct the user to the URL of a target website is considered to be a second "special technical feature."
- Claims 27-31, directed to a method of counting visits generated by a Usenet advertisement. The use of a redirecting site to count the visits is a third "special technical feature."
- Claim 44*, directed to a method for identifying a collection of images posted on Usenet. The use of author, subject and message time to identify a collection is considered to constitute a fourth "special technical feature."

Since these groups of claims do not share any of the technical features identified, a 'technical relationship' between the inventions, as defined in PCT rule 13.2 does not exist. Accordingly the international application does not relate to one invention or to a single inventive concept.

* Please note that the last four claims, following claim 40, are actually numbered 36-39, repeating the numbering of earlier claims. This report refers to these last four claims as claims 41-44.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/AU2003/001668

This Annex lists the known "A" publication level patent family members relating to the patent documents cited in the above-mentioned international search report. The Australian Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

Patent Document Cited in Search Report		Patent Family Member			
WO	0150337	AU	78927/00	US	2001054084
US	6493703	NONE			
US	5815663	NONE			
END OF ANNEX					

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ **BLACK BORDERS**
- ☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- ☒ **FADED TEXT OR DRAWING**
- ☐ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- ☐ **SKEWED/SLANTED IMAGES**
- ☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- ☐ **GRAY SCALE DOCUMENTS**
- ☐ **LINES OR MARKS ON ORIGINAL DOCUMENT**
- ☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- ☐ **OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.